

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Method and Apparatus for
Generating Random Numbers**

Inventors:

Scott A. Field

Jeffrey F. Spelman

09744784-11500

05744731-11000

1 TECHNICAL FIELD

2 The present invention relates to random number generators and, more
3 particularly, to computer-implemented random number generators that create
4 strings of random bits.

5
6 BACKGROUND

7 Random number generation is an important part of the security
8 infrastructure in many application programs and operating systems. For example,
9 random numbers are used to generate session keys and cryptographic keys for
10 encoding data that is transmitted between two locations (such as between a client
11 and a server). The use of such keys protects the integrity of the data and provides
12 for the authentication of the data and authentication of the user attempting to
13 access the data.

14 The quality of the random numbers generated is associated with the quality
15 of the security provided by the application program or operating system. A perfect
16 random number generator that produces a truly random sequence of bits is
17 considered by many to be impossible. Thus, designers attempt to create “pseudo”
18 random number generators that produce unpredictable sequences of bits in which
19 no particular bit is more likely to be generated at a given time or place in the
20 sequence than any other bit. This disclosure uses the terms “random number
21 generator” and “pseudo random number generator” interchangeably.

22
23 The quality of the random seed used by the random number generator
24 affects the quality of the random number created by the random number generator.
25 Common techniques for creating a random seed include using operating

parameters of the computer, such as time of day, date, available memory, and the like. In general, these types of parameters are regarded as sufficient for certain types of simple applications, but can lead to predictability in certain situations. For example, systems that use the computer's system time as the random seed can be predicted if the approximate system time is known, such as the time at which an email was generated. An analyst could test all possible times near the known approximate system time until the seed for the random number is discovered, thereby breaking the security of the system.

Many existing random number generation systems use data that is reset each time the computer system is reset, thereby limiting the quality of the seed data. Other random number generation systems use data that may be similar from one computer system to the next, such as time of day or date, thereby reducing the randomness of the seed data.

The system and method described herein addresses these limitations by providing a random number generator that uses random seed data that has been generated over the lifetime of the computer system.

SUMMARY

The system and methods described herein provide a computer-implemented random number generator that creates strings of random bits using entropy data that is collected over the lifetime of the computer system. The quality of this entropy data (i.e., the randomness of the data) is improved as compared to systems that use entropy data that is reset each time the computer or application program is

1 restarted. Further, the system and methods described herein collect entropy data
2 from multiple sources, thereby reducing the likelihood that two computer systems
3 will have the same entropy data. Thus, the systems and methods described herein
4 generate random numbers having an improved quality.

5 In one embodiment, entropy data is collected and stored in a nonvolatile
6 memory. The entropy data stored in the nonvolatile memory is updated with
7 newly collected entropy data. A string of random bits is generated from the
8 entropy data stored in the nonvolatile memory.

9 In a described embodiment, the entropy data is collected from multiple
10 sources within a computer system.

11 In a particular embodiment, the entropy data includes data related to a
12 processor in a computer system and data related to an operating system executing
13 on the computer system.

14 In a described implementation, the entropy data is maintained in a protected
15 portion of an operating system kernel such that the entropy data is inaccessible to
16 application programs executing on the system.

17 In one embodiment, generating a string of random bits includes hashing the
18 entropy data to generate random seed data.

19 A particular embodiment includes communicating the string of random bits
20 to an application program requesting a random number.
21
22
23
24
25

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates a block diagram of a system that collects and stores entropy data and generates strings of random bits based on the entropy data.

Fig. 2 is a flow diagram illustrating a procedure for collecting, storing, and updating entropy data.

Fig. 3 is a flow diagram illustrating a procedure for generating a string of random bits at the request of an application program.

Fig. 4 illustrates an example of a suitable operating environment in which the random number generator may be implemented.

DETAILED DESCRIPTION

The system and methods described herein provide a random number generator that creates strings of random bits using entropy data that is collected over the lifetime of the computer system. Entropy data refers to the data (such as computer system state information) used as a random seed for the random number generator. Using data that is collected over the lifetime of the computer system improves the quality of the resulting random numbers because the data is not reset each time the computer system or application program is restarted. The collected data continues to change as additional operations, functions, and application programs are executed on the computer system. As time passes, the likelihood that two different computer systems would produce the same entropy data is reduced.

Fig. 1 illustrates a block diagram of a system 100 that collects and stores entropy data and generates strings of random bits based on the entropy data. In a

particular embodiment, the system 100 is contained within a computer system. The entropy data collected includes central processing unit (CPU) data 102 and operating system data 104. CPU data 102 includes various CPU state information and operating system data 104 includes various operating system state information. Table 1 below illustrates exemplary CPU data 102 and operating system data 104.

TABLE 1

CPU Data	Operating System Data
timestamp counter	boot time
cache misses per second	time of day
branch mispredictions per second	time zone bias
CPU-specific counters	page size
	number of processors
	current cache size
	peak cache size
	I/O read operation count
	I/O write operation count
	cache read count

The CPU data 102 may vary from one CPU to the next. For example, many of the internal counters in a CPU are affected by power fluctuations, the types of operations performed by the CPU, and the clock speed at which the CPU is operating.

1 A random number generator 106 receives the collected CPU data 102 and
2 operating system data 104. The random number generator 106 stores the collected
3 data in a nonvolatile memory 108, such as a hard disk, floppy disk, flash memory
4 device or an EEPROM. Since the data is stored in the nonvolatile memory 108,
5 the data is available to the random number generator 106 after a computer system
6 restart. Thus, the data is collected and stored over the operating lifetime of the
7 computer system. The random number generator 106 is capable of processing the
8 CPU data 102 and the operating system data 104 to generate a string of random
9 bits (or bytes). Periodically, the random number generator 106 retrieves current
10 CPU data 102 and current operating system data 104. This current data is used to
11 update the data stored in nonvolatile memory 108, thereby modifying the seed data
12 used by the random number generator 106. In one embodiment, the data stored in
13 nonvolatile memory 108 is updated at regular time intervals, as controlled by a
14 timer 110. In another embodiment, the data stored in nonvolatile memory 108 is
15 updated after a particular number of requests for random numbers (e.g., after every
16 tenth request for a random number).

17 In a particular embodiment, a system device driver resides in the operating
18 system kernel and generates random numbers at the request of an application
19 program or other function accessing the device driver. This device driver is
20 responsible for collecting and maintaining entropy data as discussed herein. An
21 application programming interface (API) is provided to allow application
22 programs to request a random number. The API communicates random number
23 requests to the device driver, which generates a random number (a string of
24 random bits or bytes) based on the entropy data. In a particular implementation,
25 application programs use the RtlGenRandom() API provided by the Windows®

operating system, developed by Microsoft Corporation of Redmond, Washington. The RtlGenRandom() API communicates with the device driver via the Win32[®] application programming interface call DeviceIoControl(), which is a commonly used API call for communicating with device drivers in the Windows[®] operating system.

As discussed below, the device driver applies a hash function to the various entropy data collected. The result of the hash function is used as the random seed for the random number generator. The entropy data and the random seed data are maintained in a protected portion of the operating system kernel (i.e., a portion of the operating system kernel that is not accessible by an application program). Maintaining the entropy data and the random seed data in a protected portion of the operating system kernel prevents an application program from predicting or deriving random numbers issued to another application program on the same computer system. In a particular embodiment, the device driver manages the memory used to store the entropy data and the random seed data. In this embodiment, entropy data is maintained in the operating system kernel as well as the non-volatile Windows[®] registry.

A typical computer system has multiple processes executing simultaneously, one or more of which may require random numbers. The system and methods described herein allow the generation of multiple random numbers for use as session keys, cryptographic keys, and the like. Although particular embodiments are discussed with reference to a device driver residing in the operating system kernel that generates random numbers, it will be appreciated that any type of software component and/or firmware component can be used to implement the random number generator.

Fig. 2 is a flow diagram illustrating a procedure 200 for collecting, storing, and updating entropy data. When a system is first initialized, a device driver collects initial entropy data from the CPU and the operating system (block 202). The initial entropy data is then stored in a nonvolatile memory (block 204), such as nonvolatile memory 108 in Fig. 1. The initial entropy data is hashed to generate random seed data for the random number generator (block 206). The hashed data may be stored in a register or other storage location that is accessible to the random number generator, but inaccessible to application programs executing on the system. Any hashing algorithm that produces a long string of bits can be used to hash the entropy data. An example hashing algorithm is described in U.S. Patent 5,778,069, the disclosure of which is incorporated by reference herein. In an alternate embodiment, two or more different hashing algorithms are applied to the same set of entropy data and the results are concatenated together into a single string of bits representing the random seed data.

At block 208 in Fig. 2, the procedure 200 determines whether to update the entropy data. The entropy data may be updated at periodic intervals (e.g., every fifteen minutes) or after generating a particular number of random numbers (e.g., after every tenth random number is generated). Alternatively, an application program may specifically request an update of the entropy data. A particular implementation updates the data the first time that an application program makes a request for a random number. If the entropy data needs to be updated, then the procedure continues to block 210, where the device driver collects the current entropy data (i.e., the CPU data and the operating system data). After collecting the current entropy data, the device driver updates the data in the nonvolatile

1 memory by replacing the previous entropy data with the new entropy data (block
2 212). Alternatively, the device driver may update the data in the nonvolatile
3 memory by hashing the previous entropy data with the new entropy data (this
4 process may be referred to as "chaining"). Finally, the updated entropy data is
5 hashed to generate updated random seed data (block 214). The resulting random
6 seed data overwrites the previous seed data stored in a register or other storage
7 location. The procedure 200 then returns to block 208 to await the next update of
8 the entropy data.

9 When a system is first initialized (i.e., no entropy data is stored in the
10 nonvolatile memory), the procedure shown in Fig. 2 is executed beginning at
11 block 202. However, when a system is reset or rebooted (i.e., entropy data is
12 already stored in the nonvolatile memory), then the procedure shown in Fig. 2 is
13 executed beginning at block 210. Thus, resetting or rebooting a system does not
14 cause the deletion of any previously stored entropy data. The entropy data stored
15 in the nonvolatile memory is continually updated over the life of the system,
16 thereby providing continually changing entropy data and higher quality random
17 numbers generated from the entropy data.

18 Fig. 3 is a flow diagram illustrating a procedure 300 for generating a string
19 of random bits at the request of an application program. An application program
20 requests a random number by calling a device driver (via an API) that resides in
21 the operating system kernel (block 302). The device driver retrieves the random
22 seed data previously produced from the entropy data using a hash algorithm (block
23 304). The random seed data is retrieved, for example, from a register or other
24 storage location. Next, the device driver generates a string of random bits using
25

1 the retrieved random seed data (block 306). Finally, the device driver returns the
2 string of random bits to the requesting application program (block 308).

3 In a particular embodiment, the entropy data is hashed to produce a 640 bit
4 hash, which is the seed data for the random number generator. The random
5 number generator uses the 640 bit hash to generate a 256 byte random number,
6 which is also referred to as a "key." The 256 byte random number can be used as
7 a session key, a cryptographic key, or in any other situation requiring a random
8 number. In one implementation, the RSA RC4 stream cipher (available from RSA
9 Security of Bedford, Massachusetts) is used to generate a 256 byte random
10 number from the 640 bit hash.

11 Although particular implementations have been described above with
12 reference to specific stream ciphers, other types of ciphers can be used to generate
13 a random number from the 640 bit hash. Further, the 640 bit hash and the 256
14 byte random numbers represent an exemplary embodiment. The system and
15 methods described herein can be used with a hash (i.e., seed data) of any size to
16 generate a random number having any number of bits (or bytes).

17 Fig. 4 illustrates an example of a suitable operating environment in which
18 the random number generator may be implemented. The illustrated operating
19 environment is only one example of a suitable operating environment and is not
20 intended to suggest any limitation as to the scope of use or functionality of the
21 invention. Other well known computing systems, environments, and/or
22 configurations that may be suitable for use with the invention include, but are not
23 limited to, personal computers, server computers, hand-held or laptop devices,
24 multiprocessor systems, microprocessor-based systems, programmable consumer
25 electronics, gaming consoles, cellular telephones, network PCs, minicomputers,

1 mainframe computers, distributed computing environments that include any of the
2 above systems or devices, and the like.

3 Fig. 4 shows a general example of a computer 342 that can be used in
4 accordance with the invention. Computer 342 is shown as an example of a
5 computer that can perform the hashing and random number generation functions
6 described herein. Computer 342 includes one or more processors or processing
7 units 344, a system memory 346, and a bus 348 that couples various system
8 components including the system memory 346 to processors 344.

9 The bus 348 represents one or more of any of several types of bus
10 structures, including a memory bus or memory controller, a peripheral bus, an
11 accelerated graphics port, and a processor or local bus using any of a variety of
12 bus architectures. The system memory 346 includes read only memory (ROM)
13 350 and random access memory (RAM) 352. A basic input/output system (BIOS)
14 354, containing the basic routines that help to transfer information between
15 elements within computer 342, such as during start-up, is stored in ROM 350.
16 Computer 342 further includes a hard disk drive 356 for reading from and writing
17 to a hard disk, not shown, connected to bus 348 via a hard disk drive interface 357
18 (e.g., a SCSI, ATA, or other type of interface); a magnetic disk drive 358 for
19 reading from and writing to a removable magnetic disk 360, connected to bus 348
20 via a magnetic disk drive interface 361; and an optical disk drive 362 for reading
21 from and/or writing to a removable optical disk 364 such as a CD ROM, DVD, or
22 other optical media, connected to bus 348 via an optical drive interface 365. The
23 drives and their associated computer-readable media provide nonvolatile storage
24 of computer readable instructions, data structures, program modules and other data
25 for computer 342. Although the exemplary environment described herein employs

1 a hard disk, a removable magnetic disk 360 and a removable optical disk 364, it
2 will be appreciated by those skilled in the art that other types of computer readable
3 media which can store data that is accessible by a computer, such as magnetic
4 cassettes, flash memory cards, random access memories (RAMs), read only
5 memories (ROM), and the like, may also be used in the exemplary operating
6 environment.

7 A number of program modules may be stored on the hard disk, magnetic
8 disk 360, optical disk 364, ROM 350, or RAM 352, including an operating system
9 370, one or more application programs 372, other program modules 374, and
10 program data 376. A user may enter commands and information into computer
11 342 through input devices such as keyboard 378 and pointing device 380. Other
12 input devices (not shown) may include a microphone, joystick, game pad, satellite
13 dish, scanner, or the like. These and other input devices are connected to the
14 processing unit 344 through an interface 368 that is coupled to the system bus
15 (e.g., a serial port interface, a parallel port interface, a universal serial bus (USB)
16 interface, etc.). A monitor 384 or other type of display device is also connected to
17 the system bus 348 via an interface, such as a video adapter 386. In addition to the
18 monitor, personal computers typically include other peripheral output devices (not
19 shown) such as speakers and printers.

20 Computer 342 operates in a networked environment using logical
21 connections to one or more remote computers, such as a remote computer 388.
22 The remote computer 388 may be another personal computer, a server, a router, a
23 network PC, a peer device or other common network node, and typically includes
24 many or all of the elements described above relative to computer 342, although
25 only a memory storage device 390 has been illustrated in Fig. 4. The logical

When used in a LAN networking environment, computer 342 is connected to the local network 392 through a network interface or adapter 396. When used in a WAN networking environment, computer 342 typically includes a modem 398 or other means for establishing communications over the wide area network 394, such as the Internet. The modem 398, which may be internal or external, is connected to the system bus 348 via a serial port interface 368. In a networked environment, program modules depicted relative to the personal computer 342, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

Computer 342 typically includes at least some form of computer readable media. Computer readable media can be any available media that can be accessed by computer 342. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to,

1 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
2 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
3 tape, magnetic disk storage or other magnetic storage devices, or any other media
4 which can be used to store the desired information and which can be accessed by
5 computer 342. Communication media typically embodies computer readable
6 instructions, data structures, program modules or other data in a modulated data
7 signal such as a carrier wave or other transport mechanism and includes any
8 information delivery media. The term "modulated data signal" means a signal that
9 has one or more of its characteristics set or changed in such a manner as to encode
10 information in the signal. By way of example, and not limitation, communication
11 media includes wired media such as wired network or direct-wired connection,
12 and wireless media such as acoustic, RF, infrared and other wireless media.
13 Combinations of any of the above should also be included within the scope of
14 computer readable media.

15 The invention has been described in part in the general context of
16 computer-executable instructions, such as program modules, executed by one or
17 more computers or other devices. Generally, program modules include routines,
18 programs, objects, components, data structures, etc. that perform particular tasks
19 or implement particular abstract data types. Typically the functionality of the
20 program modules may be combined or distributed as desired in various
21 embodiments.

22 For purposes of illustration, programs and other executable program
23 components such as the operating system are illustrated herein as discrete blocks,
24 although it is recognized that such programs and components reside at various
25

